

AMENDMENTS TO CLAIMS

1. (Original) A method for mirroring data between a plurality of sites, comprising:
maintaining, at a first site of the plurality of sites, a record that identifies which
transactions that have been executed at the first site have had their redo
information replicated to the other sites of the plurality of sites;
determining a priority value associated with a transaction that is to be performed at the
first site, wherein the transaction specifies a modification to a data block;
if the priority value is a first value in a set of possible values, then committing the
transaction only after the record indicates that redo information associated with
the transaction has been replicated to the other sites of the plurality of sites; and
if the priority value is a second value in said set of possible values, then committing the
transaction even though the record does not indicate that redo information
associated with the transaction has been replicated to the other sites of the
plurality of sites.
2. (Original) The method of Claim 1, wherein the first value indicates that the transaction
should not be lost if the first site becomes inoperable.
3. (Original) The method of Claim 1, wherein the second value indicates the transaction
can be lost if the first site becomes inoperable.
4. (Original) The method of Claim 1, further comprising the step of:
determining whether all other transactions that have committed before the transaction has
committed have had their respective redo information replicated to the other sites

of the plurality of sites by comparing a commit record associated with the transaction to the record.

5. (Original) The method of Claim 1, wherein the record is a first record, and the method further comprises the step of:
maintaining, at the first site, a second record that identifies which transactions that have executed at the first site have had their redo information logged to persistent storage at the first site.
6. (Original) The method of Claim 1, further comprising the step of:
if the priority value is the second value in the set of possible values, then committing the transaction before the record indicates that the redo information generated by the transaction has been replicated to the other sites of the plurality of sites.
7. (Original) The method of Claim 5, further comprising the step of:
if the priority value is the second value in the set of possible values, then committing the transaction after the second record indicates that the redo information generated by the transaction has been stored to persistent storage at the first site.
8. (Original) The method of Claim 5, further comprising the step of:
determining which transactions that have executed at the first site have had their redo information logged to persistent storage by comparing a commit record associated with the transaction to the second record.
9. (Original) The method of Claim 1, further comprising the step of:

if a particular site of the plurality of sites becomes inoperable, then initiating recovery of the particular site after it is determined that all messages transmitted from the particular site to each other site of the plurality of sites have been received at their destination.

10. (Original) The method of Claim 1, further comprising the steps of:
at each site of the plurality of sites, determining if a data structure is to be replicated to each other site of the plurality of sites; and
replicating the data structure to each other site of the plurality of sites unless it is determined that the data structure is not to be replicated to each other site of the plurality of sites.
11. (Currently Amended) A method for storing data, comprising:
at a first site in a plurality of sites, processing a transaction;
generating in volatile memory redo information ~~that reflects for~~ the processed transaction; ~~[[and]]~~
~~delaying storing the redo information to durable storage as long as (1) if said information has not been durably stored before either~~ a data block associated with the processed transaction is not durably stored ~~[[or]]~~ and (2) the data block is not transferred to another site of the plurality of sites; and
storing the redo information to the durable storage in response to detecting that (1), then-
~~durably storing said information before either the data block is about to be~~
durably stored or (2) the data block is about to be transferred to another site of the plurality of sites.

12. (Currently Amended) The method of Claim 11, wherein the data block is a first data block, wherein the transaction is a first transaction, the redo information is [[a]] first redo information, and the method further comprises the steps of:
- at the first site, processing a second transaction;
- generating in the volatile memory second redo information ~~that reflects for~~ the processed second transaction; [[and]]
- ~~if said first information and second information has not been durably stored before either~~
- delaying storing the first redo information and the second redo information to the durable storage as long as (1) the first data block and a second data block associated with the processed second transaction [[is]] are not durably stored
- [[or]] and (2) the first data block and the second data block [[is]] are not
- transferred to another site of the plurality of sites; and
- storing, using a batch process, the first redo information and the second redo information to the durable storage in response to detecting that (1), then durably storing using a batch process said first information and said second information before either
- the first data block or the second data block is about to be durably stored or (2)
- the first data block or the second data block is about to be transferred to another site of the plurality of sites.
13. (Currently Amended) The method of Claim 12, further comprising the step of:
- determining whether the batch process has completed durably storing the first redo information and the second redo information.
14. (Original) A method for mirroring data between a plurality of sites, comprising:

maintaining, at a first site of the plurality of sites, a record that identifies which changes made to one or more data blocks stored at the first site have had associated redo information replicated to the other sites of the plurality of sites, wherein the first site implements a write-ahead logging scheme;

determining if the first site replicates, to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued; and

if the first site does not replicate, to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued, then durably storing a data block, in the one or more data blocks, associated with a transaction only after the record indicates that any write transactions that have updated the data block at the first site have had their respective redo information replicated to the other sites of the plurality of sites.

15. (Original) The method of Claim 14, wherein the record is a first record, and further comprising the steps of:

maintaining, at the first site, a second record that identifies which changes made to the one or more data blocks stored at the first site have had associated redo information logged to persistent storage at the first site; and

if the first site does replicate, to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued, then durably storing the data block after the second record indicates that any write transactions that have updated the data block at the first site have had their respective redo information logged to persistent storage at the first site.

16. (Original) The method of Claim 14, further comprising the step of:
releasing a lock associated with the data block after the first record indicates that redo
information associated with changes made to the data block has been replicated to
the other sites of the plurality of sites.
17. (Original) The method of Claim 15, wherein the first site replicates write transactions to
the other sites of the plurality of sites asynchronously to the completion of the write
transaction at the first site.
18. (Original) The method of Claim 14, further comprising the step of:
if a particular site of the plurality of sites becomes inoperable, then initiating recovery of
the particular site after it is determined that all messages transmitted from the
particular site to each other site of the plurality of sites have been received at their
destination.
19. (Original) The method of Claim 14, further comprising the steps of:
at each site of the plurality of sites, determining if a data structure is to be replicated to
each other site of the plurality of sites; and
replicating the data structure to each other site of the plurality of sites unless it is
determined that the data structure is not to be replicated to each other site of the
plurality of sites.
20. (Original) A method for mirroring data between a plurality of sites, wherein the plurality
of sites includes a first site, comprising:
at the first site, durably storing a data block prior to durably storing redo information
about changes made to the data block; and

at the first site, durably storing the redo information after the changes have been replicated to the other sites in the plurality of sites.

21. (Original) The method of Claim 20, wherein the data block is in a plurality of data blocks, wherein changes made to the plurality of data blocks are performed by transactions issued by a single process, and further comprising the step of: determining if a set of transactions issued by the single process have completed, wherein the set of transactions made the changes to the plurality of data blocks.
22. (Original) The method of Claim 20, wherein the data block is in a plurality of data blocks, wherein changes made to the plurality of data blocks are performed by transactions issued by two or more processes, and further comprising the step of: determining when the changes have been replicated to the other sites in the plurality of sites.
23. (Original) The method of Claim 20, further comprising the step of: if a particular site of the plurality of sites becomes inoperable, then initiating recovery of the particular site after it is determined that all messages transmitted from the particular site to each other site of the plurality of sites have been received at their destination.
24. (Original) The method of Claim 20, further comprising the steps of: at each site of the plurality of sites, determining if a data structure is to be replicated to each other site of the plurality of sites; and

replicating the data structure to each other site of the plurality of sites unless it is determined that the data structure is not to be replicated to each other site of the plurality of sites.

25. (Original) The method of Claim 1, wherein the record identifies which transactions that have been executed at the first site have had their redo information replicated to the other sites of the plurality of sites by identifying a portion of a redo log file, and wherein all transactions reflected in the identified portion of the redo log file have been replicated to the other sites of the plurality of sites.
26. (Original) The method of Claim 5, wherein the second record identifies which transactions that have executed at the first site have had their redo information logged to persistent storage at the first site by identifying a portion of a redo log file, and wherein all transactions reflected in the identified portion of the redo log file have been logged to persistent storage at the first site.
27. (Currently Amended) A machine-readable medium carrying one or more sequences of instructions for mirroring data between a plurality of sites, wherein the machine-readable medium is one of a volatile medium or a non-volatile medium, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:

maintaining, at a first site of the plurality of sites, a record that identifies which

transactions that have been executed at the first site have had their redo

information replicated to the other sites of the plurality of sites;

determining a priority value associated with a transaction that is to be performed at the first site, wherein the transaction specifies a modification to a data block;
if the priority value is a first value in a set of possible values, then committing the transaction only after the record indicates that redo information associated with the transaction has been replicated to the other sites of the plurality of sites; and
if the priority value is a second value in said set of possible values, then committing the transaction even though the record does not indicate that redo information associated with the transaction has been replicated to the other sites of the plurality of sites.

28. (Original) The machine-readable medium of Claim 27, wherein the first value indicates that the transaction should not be lost if the first site becomes inoperable.

29. (Original) The machine-readable medium of Claim 27, wherein the second value indicates the transaction can be lost if the first site becomes inoperable.

30. (Original) The machine-readable medium of Claim 27, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:

determining whether all other transactions that have committed before the transaction has committed have had their respective redo information replicated to the other sites of the plurality of sites by comparing a commit record associated with the transaction to the record.

31. (Original) The machine-readable medium of Claim 27, wherein the record is a first record, and wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of: maintaining, at the first site, a second record that identifies which transactions that have executed at the first site have had their redo information logged to persistent storage at the first site.
32. (Original) The machine-readable medium of Claim 27, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
if the priority value is the second value in the set of possible values, then committing the transaction before the record indicates that the redo information generated by the transaction has been replicated to the other sites of the plurality of sites.
33. (Original) The machine-readable medium of Claim 31, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
if the priority value is the second value in the set of possible values, then committing the transaction after the second record indicates that the redo information generated by the transaction has been stored to persistent storage at the first site.
34. (Original) The machine-readable medium of Claim 31, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:

determining which transactions that have executed at the first site have had their redo information logged to persistent storage by comparing a commit record associated with the transaction to the second record.

35. (Original) The machine-readable medium of Claim 27, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:

if a particular site of the plurality of sites becomes inoperable, then initiating recovery of the particular site after it is determined that all messages transmitted from the particular site to each other site of the plurality of sites have been received at their destination.

36. (Original) The machine-readable medium of Claim 27, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the steps of:

at each site of the plurality of sites, determining if a data structure is to be replicated to each other site of the plurality of sites; and

replicating the data structure to each other site of the plurality of sites unless it is determined that the data structure is not to be replicated to each other site of the plurality of sites.

37. (Currently Amended) A machine-readable medium carrying one or more sequences of instructions for storing data, wherein the machine-readable medium is one of a volatile medium or a non-volatile medium, wherein execution of the one or more sequences of

instructions by one or more processors causes the one or more processors to perform the steps of:

at a first site in a plurality of sites, processing a transaction;

generating in volatile memory redo information ~~that reflects for~~ the processed transaction; [[and]]

~~if said information has not been durably stored before either~~ delaying storing the redo information to durable storage as long as (1) a data block associated with the processed transaction is not durably stored [[or]] and (2) the data block is not transferred to another site of the plurality of sites; and
storing the redo information to the durable storage in response to detecting that (1), then-
~~durably storing said information before either~~ the data block is about to be durably stored or (2) the data block is about to be transferred to another site of the plurality of sites.

38. (Currently Amended) The machine-readable medium of Claim 37, wherein the data block is a first data block, wherein the transaction is a first transaction, the redo information is a first redo information, and wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the steps of:
- at the first site, processing a second transaction;
- generating in the volatile memory second redo information ~~that reflects for~~ the processed second transaction; [[and]]
- ~~if said first information and second information has not been durably stored before either~~ delaying storing the first redo information and the second redo information to the

durable storage as long as (1) the first data block and a second data block associated with the processed second transaction [[is]] are not durably stored [[or]] and (2) the first data block and the second data block [[is]] are not transferred to another site of the plurality of sites; and
storing, using a batch process, the first redo information and the second redo information to the durable storage in response to detecting that (1), then durably storing using a batch process said first information and said second information before either the first data block or the second data block is about to be durably stored or (2) the first data block or the second data block is about to be transferred to another site of the plurality of sites.

39. (Currently Amended) The machine-readable medium of Claim 38, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
determining whether the batch process has completed durably storing the first redo information and the second redo information.
40. (Currently Amended) A machine-readable medium carrying one or more sequences of instructions for mirroring data between a plurality of sites, wherein the machine-readable medium is one of a volatile medium or a non-volatile medium, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of:
maintaining, at a first site of the plurality of sites, a record that identifies which changes made to one or more data blocks stored at the first site have had associated redo

information replicated to the other sites of the plurality of sites, wherein the first site implements a write-ahead logging scheme;

determining if the first site replicates, to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued; and

if the first site does not replicate, to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued, then durably storing a data block, in the one or more data blocks, associated with a transaction only after the record indicates that any write transactions that have updated the data block at the first site have had their respective redo information replicated to the other sites of the plurality of sites.

41. (Original) The machine-readable medium of Claim 40, wherein the record is a first record, and wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the steps of:
- maintaining, at the first site, a second record that identifies which changes made to the one or more data blocks stored at the first site have had associated redo information logged to persistent storage at the first site; and
- if the first site does replicate , to the other sites of the plurality of sites, write transactions that are executed at the first site in the order in which the write transactions were issued, then durably storing the data block after the second record indicates that any write transactions that have updated the data block at the first site have had their respective redo information logged to persistent storage at the first site.

42. (Original) The machine-readable medium of Claim 40, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
- releasing a lock associated with the data block after the first record indicates that redo information associated with changes made to the data block has been replicated to the other sites of the plurality of sites.
43. (Original) The machine-readable medium of Claim 41, wherein the first site replicates write transactions to the other sites of the plurality of sites asynchronously to the completion of the write transaction at the first site.
44. (Original) The machine-readable medium of Claim 40, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
- if a particular site of the plurality of sites becomes inoperable, then initiating recovery of the particular site after it is determined that all messages transmitted from the particular site to each other site of the plurality of sites have been received at their destination.
45. (Original) The machine-readable medium of Claim 40, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the steps of:
- at each site of the plurality of sites, determining if a data structure is to be replicated to each other site of the plurality of sites; and

replicating the data structure to each other site of the plurality of sites unless it is determined that the data structure is not to be replicated to each other site of the plurality of sites.

46. (Currently Amended) A machine-readable medium carrying one or more sequences of instructions for mirroring data between a plurality of sites, wherein the plurality of sites includes a first site, wherein the machine-readable medium is one of a volatile medium or a non-volatile medium, wherein execution of the one or more sequences of instructions by one or more processors causes the one or more processors to perform the steps of: at the first site, durably storing a data block prior to durably storing redo information about changes made to the data block; and at the first site, durably storing the redo information after the changes have been replicated to the other sites in the plurality of sites.
47. (Original) The machine-readable medium of Claim 46, wherein the data block is in a plurality of data blocks, wherein changes made to the plurality of data blocks are performed by transactions issued by a single process, and wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of: determining if a set of transactions issued by the single process have completed, wherein the set of transactions made the changes to the plurality of data blocks.
48. (Original) The machine-readable medium of Claim 46, wherein the data block is in a plurality of data blocks, wherein changes made to the plurality of data blocks are performed by transactions issued by two or more processes, and wherein execution of the

one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:
determining when the changes have been replicated to the other sites in the plurality of sites.

49. (Original) The machine-readable medium of Claim 46, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the step of:

if a particular site of the plurality of sites becomes inoperable, then initiating recovery of the particular site after it is determined that all messages transmitted from the particular site to each other site of the plurality of sites have been received at their destination.

50. (Original) The machine-readable medium of Claim 46, wherein execution of the one or more sequences of instructions by the one or more processors causes the one or more processors to further perform the steps of:

at each site of the plurality of sites, determining if a data structure is to be replicated to each other site of the plurality of sites; and
replicating the data structure to each other site of the plurality of sites unless it is determined that the data structure is not to be replicated to each other site of the plurality of sites.

51. (Original) The machine-readable medium of Claim 27, wherein the record identifies which transactions that have been executed at the first site have had their redo information replicated to the other sites of the plurality of sites by identifying a portion of

a redo log file, and wherein all transactions reflected in the identified portion of the redo log file have been replicated to the other sites of the plurality of sites.

52. (Original) The machine-readable medium of Claim 31, wherein the second record identifies which transactions that have executed at the first site have had their redo information logged to persistent storage at the first site by identifying a portion of a redo log file, and wherein all transactions reflected in the identified portion of the redo log file have been logged to persistent storage at the first site.
53. (Original) The method of Claim 14, wherein the record identifies which changes are made to the one or more data blocks stored at the first site have had associated redo information replicated to the other sites of the plurality of sites by identifying a portion of a redo log file, and wherein all changes in the identified portion of the redo log file have been replicated to the other sites of the plurality of sites.
54. (Original) The method of Claim 15, wherein the second record identifies which changes are made to the one or more data blocks stored at the first site have had associated redo information logged to persistent storage by identifying a portion of a redo log file, and wherein all changes in the identified portion of the redo log file have been logged to persistent storage.
55. (Original) The machine-readable medium of Claim 40, wherein the record identifies which changes are made to the one or more data blocks stored at the first site have had associated redo information replicated to the other sites of the plurality of sites by identifying a portion of a redo log file, and wherein all changes in the identified portion of the redo log file have been replicated to the other sites of the plurality of sites.

56. (Original) The machine-readable medium of Claim 41, wherein the second record identifies which changes are made to the one or more data blocks stored at the first site have had associated redo information logged to persistent storage by identifying a portion of a redo log file, and wherein all changes in the identified portion of the redo log file have been logged to persistent storage.